

Claims

- [c1] A device comprising:
- a port to receive one or more data streams, each data stream including one or more data frames;
 - a task scheduler coupled to the port, the task scheduler to generate a task identifier for every data frame received;
 - a first queue coupled to the task scheduler to hold task identifiers of a first priority type;
 - a second queue couple to the task scheduler to hold task identifiers of a second priority type, the second priority type different than the first priority type;
 - a switch coupled to the first and second queues, the switch configured to retrieve task identifiers from the first queue and the second queue in a fair manner; and
 - a third queue coupled to the switch, the third queue to hold task identifiers retrieved by the switch.
- [c2] The device of claim 1 further comprising:
- a classifier communicatively coupled to the port to assign a priority type to every data frame received.
- [c3] The device of claim 1 further comprising:
- a look-up table store communicatively coupled to the port, the look-up table store to store conversions between priority types and data frame types.
- [c4] The device of claim 3 wherein the conversions between priority types and data frame types are pre-assigned.
- [c5] The device of claim 3 wherein the conversions between priority types and data frame types are dynamically configured.
- [c6] The device of claim 1 further comprising:
- a task router coupled to receive task identifiers from the task scheduler and the task identifiers in either the first or second queue.
- [c7] The device of claim 1 wherein the task router is configured to monitor the first queue for an overflow condition and, if an overflow condition is detected,

reassign data frame priority types to prevent overflow of the first queue.

- [c8] The device of claim 1 wherein the switch is configured to retrieve task identifiers from the first and second queues in a fair and weighted manner according to priority types of the task identifiers.
- [c9] The device of claim 8 wherein task identifiers with a higher priority type are weighted more heavily than task identifiers with a lower priority type.
- [c10] The device of claim 1 wherein the third queue is a shared execution queue from which one or more processing units retrieve task identifiers to process.
- [c11] A method comprising:
receiving one or more data streams, each data stream including one or more data frames of one or more data frame types;
determining a task priority level for each data frame received;
routing each data frame to one of one or more storage queues based on the task priority level of each data frame; and
retrieving the data frames from the one or more storage queues during a task retrieval cycle according to a fair and weighted processing scheme based on task priority level.
- [c12] The method of claim 11 wherein the task priority level is determined from the data frame type.
- [c13] The method of claim 11 wherein the task priority level corresponding to a particular data frame type is pre-configured.
- [c14] The method of claim 11 wherein each storage queue stores data frames of a different task priority level than the other storage queues.
- [c15] The method of claim 11 wherein each data frame type corresponds to particular processing time requirements.
- [c16] The method of claim 15 wherein, according to the weighted processing scheme, data frames of approximately equal total processing time restrictions are retrieved from each storage queue in a task retrieval cycle.

- [c17] The method of claim 11 wherein, if an overflow condition is detected in a first storage queue, one or more data frame types are reassigned to different priority levels to prevent further data frames from being stored in the first storage queue.
- [c18] The method of claim 11 further comprising:
placing the retrieved data frames into an execution queue to be processed.
- [c19] A method comprising:
receiving one or more data streams, each data stream including one or more data frames of one or more data frame types;
determining a task priority level for each data frame received;
assigning a unique task identifier to each received data frame;
storing each task identifier to one of multiple storage queues according to the task priority level of the corresponding data frame; and
retrieving task identifiers from the one or more storage queues during a task retrieval cycle according to a weighted processing scheme based on task priority levels.
- [c20] The method of claim 19 wherein the task priority level is determined from the data frame type.
- [c21] The method of claim 19 wherein each data frame type corresponds to particular processing time requirements.
- [c22] The method of claim 21 wherein, according to the weighted processing scheme, task identifiers corresponding to data frames of approximately equal total processing time requirements are retrieved from each storage queue in a task retrieval cycle.
- [c23] The method of claim 19 further comprising:
placing the retrieved task identifiers into an execution queue to be processed.
- [c24] A machine-readable medium having one or more instructions for scheduling processing tasks, which when executed by a processor, causes the processor to perform operations comprising:

receiving one or more data streams, each data stream including one or more data frames of one or more data frame types;
determining the task priority level for each of the data frames received;
routing each data frame to one of one or more storage queues based on the task priority level of each data frame; and
retrieving the data frames from the one or more storage queues during a task retrieval cycle according to a fair and weighted processing scheme based on task priority levels.

- [c25] The machine-readable medium of claim 24 wherein the task priority level is determined from the data frame type.
- [c26] The machine-readable medium of claim 24 wherein each storage queue stores data frames of a different task priority level than the other storage queues.
- [c27] The machine-readable medium of claim 24 wherein each data frame type corresponds to particular processing time requirements.
- [c28] The machine-readable medium of claim 27 wherein, according to the weighted processing scheme, data frames of approximately equal total processing time are retrieved from each storage queue in a task retrieval cycle.
- [c29] The machine-readable medium of claim 24 wherein, if an overflow condition is detected in a first storage queue, one or more data frame types are reassigned to different priority levels to prevent further data frames from being stored in the first storage queue.
- [c30] The machine-readable medium of claim 24 further comprising:
placing the retrieved data frames into an execution queue to be processed.